openstack.

November 2017

OpenStack Queens Roadmap

OpenStack Summit Sydney

openstack.

Table of Contents

openstack.

Roadmap Changes

openstack.

Previous Roadmap Model

1. PTL survey: Asked to select from and rank UX, Scalability, Modularity, Resiliency, Security, Interoperability in the upcoming release

2. Roadmap team would aggregate results and say, “The Newton release will focus on User Experience and Modularity.”

openstack.

Previous Roadmap Model

Dev Feedback:

“Always working on UX *and* security *and* manageability…ranking these doesn’t make sense.”

“We have a feature users have been asking for; I don’t know how to call it out in this survey.”

openstack.

Previous Roadmap Model

User Feedback:

“I’m interested in what’s coming for containers.”

“I’m interested in what’s coming for edge.”

What “level” should the Roadmap fly at?

30,000 foot view vs 10,000 vs 100

Users want actionable information to help them understand what’s coming in the next OpenStack release

We need your help to determine that level—see the last slide on how to help the Roadmap team!

openstack.

The Queens Roadmap

Sourced from notes from the Denver Project Teams Gathering, Blueprints, and Specs

Goal was to give a preview of 10 features that will make a difference for end users. Not an exhaustive list or a guarantee of feature delivery (Queens doesn’t come until February 2018).

openstack.

Former Roadmap Category

Potential Use Case

Queens Feature

What is it:

Brief feature synopsis

Why it matters:

How this feature will make a difference for users

openstack.

Manageability

Faster path to the current release

Fast Forward Upgrades

What it is:

Allows users to quickly move through upgrades to get from release N to N+X where X is greater than 1 and equal or less than 3, for example, going from Newton to Queens. This feature was originally called “skip level upgrades,” but you aren’t technically skipping, because there are certain things that still need to be done at each release level, like database migrations, but this part of the upgrade is limited to a single host to make it move forward as quickly as possible. This effort is being led by the TripleO project team in collaboration with other projects.

Why it matters:

OpenStack is on a 6 month release cycle, and for some operators this doesn’t fit with their organization’s culture or workflow, and they find themselves farther behind the current release than they’d like to be, or on an EOLed release. In the April 2017 User Survey, the average respondent was on N-2 releases, so this feature is addressing a huge community need. This work is just in its beginning stages, but we’ll see it appear in Queens and beyond.

openstack.

Manageability

Register and document policy in code

What it is:

Currently, the role-based access controls for the APIs of OpenStack services are usually in the project source, and they document the default policy values. The Technical Committee has made it a goal for the Queens release for projects to register and document those RBAC policies in code.

Why it matters:

This will improve the experience and manageability for both developers and operators in a couple of ways including having documentation describing each policy generated and available to assist operators and having policy files that only contain the policies for their deployment, making auditing easier. Overall, this goal means better communication about service policies and the ability to set more granular defaults in RBAC policies.

openstack.

Manageability

Containers

Edge

OpenStack-Helm

What it is:

Helm is a package manager for K8S that lets you define, install and upgrade applications on K8S. It uses “helm charts”, which is a collection of files that describe a related set of K8S resources. OpenStack-Helm is the collection of charts and tools that let you use Helm to manage the lifecycle of OpenStack on top of K8S.

Why it matters:

This is a full lifecycle management solution that lets users easily deploy and manage individual components or a full stack. Because each service has its own chart, you don’t have to run all the OpenStack services, and you can manage them independently. You can also use a parent-child set up for your charts and manage just the parent. In a Boston Summit demo, the OpenStack-Helm team had 3 bare metal hosts provisioned with K8S. They installed OpenStack mitaka in 6 minutes. After that, they upgraded to Newton in 3. OpenStack-Helm also keeps dependencies in the charts themselves, so control is pushed to the edge, rather than having a centralized controller.

openstack.

Manageability

Resiliency

Ironic Rescue Mode

What is it:

Nova rescue lets you repair an instance when something won't boot correctly or you lose an SSH key, etc. Currently, Ironic does not implement this, which means end users don't have too many options for troubleshooting misconfigured nodes for bare metal. The work to have Ironic rescue mode come to life was started in Newton, and is aimed to be finished in Queens.

Why it matters:

Nova rescue can be a critical lifeline; extending that same benefit to Ironic. In the most recent user survey report, there was a significant growth in production usage of Ironic from 9% to 20% between April 2017 and November 2017. With Ironic being run in production by more users, this is a great feature to have for the next release.

openstack.

Manageability

User Experience

User-generated templates

HOT Drag and Drop Interface

What is it:

Horizon will be adding a plugin for Heat Orchestration Templates, also known as HOT templates, to be created through a drag-and-drop UI on the Horizon dashboard.

Check out a demo: https://www.youtube.com/watch?v=Qb_biqCOxH0

Why it matters:

Currently HOT templates are plaintext files. This is a super user friendly, efficient way to orchestrate resources!

openstack.

Scalability

AI

Machine Learning

Compute-intensive Apps

Support for vGPUs

What it is:

This comes from Nova, who will be adding vGPU support so cloud admins will be able to define flavors that request vGPU resources and specify resolutions for vGPUs, and end users will be able to boot vms which have vGPUs.

Why it matters:

vGPUs, virtual graphics processing units, are incredibly powerful for certain types of high-performance workloads. Given the name, you do see them in graphics-intensive workloads, but you also see them used in a lot of scientific workloads, machine learning, AI. This change opens up an array of possibilities for end users who need that.

openstack.

Manageability

Resiliency

Clustered Applications

Cinder Multi-Attach

What it is:

The Cinder multi-attach feature means you can attach the same, single Cinder volume to multiple vm instances.

Why it matters:

Volume multi-attach is one of the most highly-requested features in cloud environments. An obvious benefit is that you can have two nodes accessing the same volume, so if one goes down, the other can take over and has access to the data.

openstack.

Manageability

Resiliency

High Availability

OpenStack Masakari

What is it:

This is another new project. Masakari helps OpenStack clouds achieve high availability from various vm failure events and automates the rescue mechanism.

Why it matters:

Good engineering is planning for failures. And with the Masakari project, it’s one more tool to help users to achieve HA.

openstack.

Manageability

User Experience

S3 API in Swift Repo

What it is:

For Queens, the Swift team is planning to integrate an S3 API compatibility layer back into the Swift repo. This means that both the Swift API and the S3 API are managed in the same code base.

Why it matters:

This feature is all about providing an easier experience for end users who might want to have their app that talks to S3 talk to Swift, and vice versa. Operators are able to give their users access to more existing applications and tools with this feature.

openstack.

Manageability

Modularity

Containers

OpenStack LOCI

What is it:

A new project, LOCI is a project that makes Open Container Initiative-compatible images of OpenStack services.

Why it matters:

LOCI's goal is to build compact but complete OCI-compatible containers that are independent of any particular deployment or orchestration technology. These containers can be dropped into heavy-weight deployment tools like OpenStack-Helm, or used individually to deliver standalone services like Cinder block storage.

openstack.

Queens, coming February 2018

openstack.

Future Roadmaps

Do you use the Roadmap?
What does an ideal Roadmap look like?

Share your feedback: anne@openstack.org

openstack.

Queens Roadmap
OpenStack Summit Sydney


openstack


@OpenStack


openstack


OpenStackFoundation